

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE January 31, 1996		3. REPORT TYPE AND DATES COVERED Final Report, 12/01/92 - 11/30/95	
4. TITLE AND SUBTITLE Probabilistic Error Checkers				5. FUNDING NUMBERS F49620-93-1-0082	
6. AUTHOR(S) Dr. Peter B. Danzig					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Southern California, Los Angeles, CA 90089-0781				AFOSR-TR-96 0330	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Office of Scientific Research 110 Duncan Avenue, Suite B115 Bolling AFB, DC 20332				10. SPONSORING/MONITORING AGENCY REPORT NUMBER 19960624 240	
11. SUPPLEMENTARY NOTES					
12a. DISTRIBUTION/AVAILABILITY STATEMENT Distribution Unlimited				12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) Abstract This final report summarizes the progress made in during the areas of research supported by AFOSR. During the contract, <ul style="list-style-type: none"> (a) the PI implemented and deployed two probabilistic error checkers, one for the domain name system and the other in an application that replicates FTP-servers; (b) the PI and his student Sugih Jamin developed a measurement-based admission control algorithm for real-time traffic over integrated services packet networks. This admission control algorithm can achieve double the performance of the competition's algorithm for voice, video, and self-similar traffic. (c) the PI, co-funded by ARPA, developed the Harvest hierarchical object cache for the Internet. The Harvest cache reduces network traffic from routine requests, improves object availability, and isolates the network from accidentally looping requests. It is in use at several hundred sites across the Internet. During the grant period, the PI won the NSF National Young Investigator (NYI) award and a 1993 USC innovative teaching award. His PhD student, Sugih Jamin, received the 1995 ACM SIGCOMM best student paper award for research co-funded by this AFOSR grant.					
14. SUBJECT TERMS computer networks, software engineering, network protocols, multi-media networks, network admission control, error detection				15. NUMBER OF PAGES 7	
				16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL		

Final Report— Probabilistic Error Checkers: F49620-93-1-0082

Peter B. Danzig

Computer Science Department
University of Southern California
Los Angeles, California 90089-0781
danzig@usc.edu

Abstract

This final report summarizes the PI's research in three related areas of research: distributed system error detection, reliable internet software systems, and measurement-based admission control for guaranteed quality of service networks.

- the PI implemented, deployed, and maintained a probabilistic error checker (PEC) for the Domain Name System (DNS). He then re-applied the PEC concept to the design of a system to massively replicate databases around the Internet.
- The PI devised a measurement-based call admission control algorithm for real-time traffic over integrated services packet networks. This admission control algorithm achieve doubles the performance of the competition's algorithm for voice, video, and self-similar traffic.
- The PI, in collaboration with other agencies, applied his research to the design of the Harvest hierarchical Internet object cache. Besides reducing network traffic from routine requests, it improves object availability and isolates the network from accidentally looping requests.

Objectives

This project developed techniques to build and debug robust, wide-area-network computer systems. The research addressed how to

- uncover hidden implementation errors in deployed software,
- isolate the effects of faulty nodes on wide-area-network performance,
- most efficiently use limited network bandwidth.

The PI developed tools to diagnose and resolve performance and design problems of heterogeneous, autonomously managed, distributed systems, focusing attention on key Internet components, including the Domain Name System (DNS) and the WWW, gopher, and FTP services.

The research was based on the premise that distributed systems fail in stereotypical ways: they deadlock, they fail half-way through an operation, their networks partition or exhibit asymmetric communication paths, their logic gets caught in loops. The PI tested that hypothesis that the way to make a robust, heterogeneous distributed system is to program its elements so that the components themselves exhibit this behavior right from the start. Applying his ideas to software prototypes, the PI demonstrated that when components of a distributed system occasionally mimic their own failure during normal service, that the system can be designed to diagnose implementation and design mistakes in other system components.

Summary of Research Effort

The PI his DNS checker and probabilistic error checker software and explored the probabilistic checker technique in the Harvest wide-area-network data replicator. He integrated the ideas behind checker into the Harvest Web cache. ISI is proposing to regain administrative control of the Internet's root name servers and deploy the PI's DNS checker software.

In a separate research thrust, he developed network conversation admission control algorithm for multi-media network applications.

Probabilistic Error Checkers

The first goal of this project was to create diagnostic software that identifies broken components of distributed services under normal service and occasionally

stress tests particular components to explore their behavior under duress. As an example of such software, the PI implemented a diagnostic tool for name server traffic. This checker program analyzed the interaction between a name server and its possibly mis-behaving clients.

The PI's probabilistic error checker for DNS demonstrates that a simple module can identify misbehaving components of a functioning distributed system by feigning death and do so without breaking the system. In contrast to the full DNS checker the probabilistic checker is significantly smaller, simpler, less memory demanding, but still capable of diagnosing broken components [2].

This software is deployed on a secondary of the ".us" domain and can be queried from <http://excalibur.usc.edu/research/checker>. The software and instruction manuals for the DNS checker software are available from <ftp://catarina.usc.edu/pub/danzig/home.html>.

Hierarchical Object Cache

The explosion of interest in the Internet information "Mosaic" has taxed the Internet's WWW, gopher, and FTP servers. Jointly with funding from Hughes Information Systems (via a NASA subcontract) and directly from ARPA, the PI built the Harvest Object Cache, which is architecturally similar to DNS, but designed to avoid the flaws that uncovered in DNS [1]. The cache has lead to an international collaboration of several hundred users <http://www.nlanr.net/Cache>

When a referenced object causes a cache miss, the cache estimates its network distance to the object's home node and checks to see if the object would hit the cache's immediate parents and siblings. If the object's home node is closer than in any sibling or parent cache for which the object is a hit, the cache fetches the object directly. Otherwise, if the object is a miss in all siblings and parents, then the cache fetches the object through the closest parent cache. By closest cache, we mean the parent cache with shortest round trip time. An example cache topology is shown in Figure 1.

Real-Time, Multi-media Networks

The PI developed a measurement-based, conversation admission control algorithm for real-time, multi-media networks [4, 3]. The applies to integrated packet networks, whether IP or ATM. To provide real-time, multi-media services, the network must assign and reserve re-

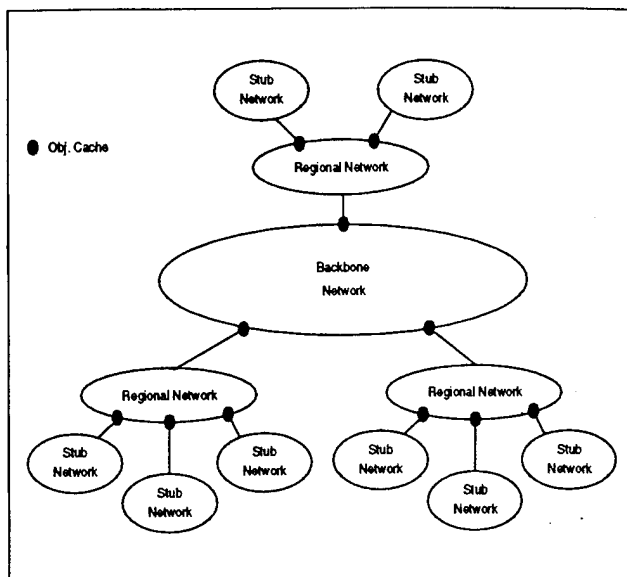


Figure 1: Hierarchical object cache for the Internet.

sources to users. While current IP and ATM networks do not yet have this infrastructure in place, they eventually will.

Existing admission control algorithms require that network users accurately specify their quality of service (QOS) requirements. The admission control algorithm then translates this QOS specification into bandwidth, buffer space, and switching priority reservations at each network switch and link.

The PI is currently developing a measurement-based admission control algorithm. The user still states his quality of service requirements, but the network measures the network's total resource usage and attempts to over-book the reservations, to increase network utilization. The admission control algorithm works with two classes of service: guaranteed service and predictive service. Guaranteed and predictive service, like the names imply, are comparable to hard and soft-real time systems. Guaranteed service reserves more resources than predictive service, but the network always meets requested QOS.

The goal of our admission control research, like our cache and replicator research, is to more efficiently use network bandwidth.

Personnel

The following students were funded by this AFOSR project:

- Steve Miller, who wrote the DNS checker,

- Charles Chan, who wrote the probabilistic error checker for DNS and who created the small-memory version of checker, is now employed at Motorola.
- Katia Obraczka, who recently completed her Ph.D. dissertation on "Massively Replicating Services in Wide Area Internetworks", now works at ISI.
- Erhyuan Tsai debugged our *mirrord* replication tool using a probabilistic error checker. Now works at Netscape.
- Anawat Chankhunthod and Chuck Neerdales worked on the hierarchical object cache. Chuck now works for Netscape. Anawat is still a PhD student.
- Sugih Jamin passed his Ph.D. thesis defense on "A Measurement-based Admission Control Algorithm for Integrated Services Packet Networks". He won best student paper prize at SIGCOMM 95 and is now interview for faculty and research lab positions.

- [3] Sugih Jamin, Peter B. Danzig, Scott Shenker, and Lixia Zhang. A measurement-based admission control algorithm for integrated services packet networks. *To appear: ACM Transactions on Networking*, 1996.
- [4] Sugih Jamin, Peter B. Danzig, Scott Shenker, and Lixia Zhang. A measurement-based admission control algorithm for integrated services packet networks. *SIGCOMM 95*, August 1995.

Awards

- received best-student paper award 1995 ACM SIGCOMM
- served on the 1993, 1995, & 1996 ACM SIGCOMM program committee
- served on the 1996 IEEE INFOCOMM program committee
- served on an NSF Networking Research Panel, December 1994 & March 1996
- served on the 1993 & 1996 ACM SIGMETRICS program committees
- received the NSF NYI award, 1994.
- Served as associate editor, Journal of Internet-working

References

- [1] Anawat Chankhunthod, Peter B. Danzig, Chuck Neerdaels, Mike Schwartz, and Kurt J. Worrell. A hierarchical internet object cache. *USENIX 1996*, January, 1996.
- [2] Peter B. Danzig, Katia Obraczka, and Anant Kumar. An analysis of wide-area name server traffic: A study of the domain name system. *Proc. of the ACM SIGCOMM '92, Baltimore, Maryland*, January 1992.

Peter B. Danzig

Address:	Computer Science Department	Office:	+1 213 740 4780
	University of Southern California	Home:	+1 310 559 8530
	941 West 37th Place	Fax:	+1 310 559 7053
	Los Angeles, California 90089-0781	Email:	danzig@usc.edu

RESEARCH INTERESTS: Computer networks and Internet information systems.

EDUCATION

- Ph.D. in Computer Science, University of California at Berkeley, December, 1989.
Dissertation Advisor: Domenic Ferrari.
- B.S. in Applied Physics, with highest honors, University of California at Davis, June 1982.

ACADEMIC EXPERIENCE

- Assistant Professor, Computer Science Department, University of Southern California. Since 1990.

AWARDS and HONORS:

- National Young Investigator (NYI), National Science Foundation, Sept 1994-1999.
- U.S.C. Innovative Teaching Award for creating CS402L 'Multimedia Operating Systems and Networks Laboratory', 1993. Awarded by the USC Innovative Teaching Committee.
- U.C. Berkeley Demetri Angelakos Service Award for assistance to fellow computer science graduate students, 1987.
- American Electronic Association Faculty Development Fellowship 1985-1989.
- Graduated with highest honors, U.C. Davis, June 1982. Distinction awarded to the best one percent of undergraduates.
- President's Undergraduate Fellowship, U.C. Davis, 1981-1982.

PROGRAM COMMITTEES and REFEREEING:

- Associate Editor: Journal of Internetworking: Research and Experience, August 1994-present.
- Program Committees: IEEE INFOCOMM 1996, ACM SIGCOMM (1996, 1995, and 1993), ACM SIGMETRICS (1993 and 1996).
- Panell Member for NSF Networking Research, December 1994, March 1996. Also referee proposals for NSF Networking research and NSF Operating Systems & Systems Software programs.
- Frequent referee for ACM SIGCOMM, ACM SIGMETRICS, ACM TOCS, ACM/IEEE Transactions on Networking, IEEE TSE (Partial List).
- University Library Committee 1992-1993. University Innovative Teaching Committee 1992-1993.
- Computer Science Department: Graduate Student Assistantship Committee 1993-1995, PhD Advising Committee 1995-

FUNDING:

- NSF, National Young Investigator, \$275,000, 9/94-9/99.
- NIH, Human Brain Project. \$4,500,000 in collaboration with Michael Arbib and 7 other USC faculty. 9/94-8/99.
- Hughes Aircraft "Object Caching and Bulk Transfer Protocols for EOSDIS, \$95,000, 2/94-7/94
- ARPA "Resource Discovery", \$432,000 subcontract of a \$2,000,000 contract, 9/93-8/96
- USC Innovative Teaching Award. \$6,600. 7/93-6/94.
- NSF Small Scale Infrastructure Grant. \$1,076,000, 1/93-12/97 (In collaboration with Estrin, Ghandeharizadeh, Horowitz, Hull, and McLeod).
- U.S Army, SBIR consultant to ACSC, \$20,000, 1/94-12/94
- Air Force Office of Scientific Research. \$243,000, 12/92-11/95, "Probabilistic Error Checkers".
- Powell Foundation Instructional Equipment Grant. \$21,000, 9/92-8/93, "Undergraduate Multimedia OS and Network Lab".
- Powell Foundation Instructional Equipment Grant. \$21,000, 9/91-5/92, "Graduate Student Computing".
- USC Innovative Teaching Award. \$6,000. 6/91-12/91.
- Powell Foundation Research Equipment Grant. \$75,000. 2/90-1/91.

REFEREED JOURNAL PUBLICATIONS

- [95a] C. Mic Bowman, Peter B. Danzig, Darren Hardy, Udi Manber, Michael F Schwartz. "The Harvest Information and Access System." *Computer Networks and ISDN Systems*. Vol 28, P 119-125, 1995.
- [94a] Shih-Hao Li and Peter B. Danzig. "Boolean Similarity Measures". *Submitted to IEEE Transactions on Knowledge Engineering*.
- [94b] Jong Suk Ahn and Peter B. Danzig. "Hybrid Technique for Simulating High Bandwidth-Delay Computer Networks." *Submitted to ACM Transactions on Networking*, April 1995.
- [94c] Peter B. Danzig. "Multicast and buffer overflow." *IEEE Transactions on Software Engineering*, Volume 20, Number 1, January 1994, 1-12.
- [93a] C. Mic Bowman, Peter B. Danzig, and Michael F. Schwartz. "Research Problems for Scalable Internet Resource Discovery". *Communications of the ACM*, Volume 37, Number 8, p 98-107, August 1994.
- [93b] Peter B. Danzig, Katia Obraczka, and Shih-Hao Li. "Internet Resource Discovery Services". *IEEE Computer*, Volume 26, Number 9, p. 8-22, September 1993.
- [92a] Peter B. Danzig, Shih-Hao Li, and Katia Obraczka, "Distributed Indexing of Autonomous Internet Services". *Computer Systems*, Volume 5, Number 4, 1992.
- [91a] Peter B. Danzig, Sugih Jamin, Ramon Caceres, Danny Mitzel, and Deborah Estrin. "An artificial workload model of a TCP/IP internetwork". *Journal of Internetworking: Research and Experience*, Vol. 3, p. 1-26, 1992.
- [91b] Peter B. Danzig and Steven Melvin "Microsecond Resolution Timing with Sun Workstations". *Behavior Research Methods, Instruments, & Computers*. Volume 23:1, p. 61-65, January 1991.
- [91c] Peter B. Danzig. "A Cooperative Game with Applications to Computer Networks" *Information Processing Letters*, Vol 38, Number 6, p. 283-289, June 1991.
- [89a] Lawrence Rowe, Peter B. Danzig, and Wilson Choi. "A Visual Shell Interface to a Database." *Software Practice and Experience*, Volume 18, p. 108-117, June 1989.

REFEREED CONFERENCE PUBLICATIONS

- [96a] Anawat Chankhunthod, Peter B. Danzig, Chuck Neerdaeles, Michael F. Schwartz, and Kurt Worrell. A Hierarchical Internet Object Cache. *1996 Usenix Technical Conference*, January 1996.
- [95b] Sugih Jamin, Peter B. Danzig, Scott Shenker, Lixia Zhang. A measurement-based Admission Control Algorithm for Integrated Services Packet Networks. To appear 1995 ACM SIGCOMM, August 1995. Recommended for publication in ACM Transactions on Networking.
- [95c] Peter B. Danzig, Zhen Liu, Limin Yan, Jong-suk Ahn. TCP Vegas: Experimental Evaluation of Packet Network Flow Control Algorithms by Live Emulation. To Appear: ACM SIGCOMM 1995, August 1995.
- [95d] Shih-Hao Li and Peter B. Danzig. "Vocabulary Problem in Internet Resource Discovery systems". To Appear: 2nd International Workshop on Next Generation Information Technologies and Systems. NGITS '95, June 1995.
- [94e] Mic Bowman, Peter Danzig, Darren R. Hardy, Udi Manber, and Michael Schwartz. "The Harvest Resource Discovery and Access System". Proceedings of the Second International World Wide Web Conference, p. 763-771, October 1994.
- [93c] Peter B. Danzig, Rick Hall, Michael F. Schwartz. "A Case for Caching File Objects Inside Internetworks." ACM SIGCOMM 93, p.239-248, September 1993.
- [93d] C. Mic Bowman, Peter B. Danzig, and Michael F. Schwartz. "Research Problems for Scalable Internet Resource Discovery" INET 93, p. DFB1-DFB10, August 1993.
- [92b] Peter B. Danzig, Anant Kumar, and Katia Obraczka. "An analysis of wide-area DNS name server traffic." ACM SIGCOMM 92, p. 281-292, August 1992.
- [91d] Ramon Caceres, Peter B. Danzig, Sugih Jamin, Danny Mitzel. "Characteristics of individual application conversations in wide area TCP/IP internetworks". 1991 ACM SIGCOMM Conference, p. 101-112, September 1991.
- [91e] Peter Danzig, Jong Suk Ahn, John Noll, and Katia Obraczka. "Distributed Indexing: A scalable mechanism for distributed information retrieval. ACM SIGIR Conference, p. 220-229, October 1991.
- [91f] Peter B. Danzig. "An Analytical Model of Operating System Protocol Processing Including Effects of Multiprogramming". ACM Sigmetrics Conference, p. 11-20, May 1991.
- [89b] Peter B. Danzig. "Finite Buffers and Fast Multicast." 1989 ACM SIGMETRICS and Performance 89 Conference, p. 108-117, May 23-26, 1989.
- [87a] Gaetano Borriello, Andrew Cherenon, Peter B. Danzig, and Michael Nelson. "Special or General-Purpose Hardware for Prolog: A Case Study." ACM ASPLOS-2, Palo Alto, October 5-8, p. 136-145, 1987.

TECHNICAL REPORTS

- [93e] Peter B. Danzig, Dante Delucia, Katia Obraczka, and Kittinon Wangpattanamongkol. "Massively Replicating Services in Autonomously Managed Wide-Area Internetworks", USC TR 93-541, March 1993.
- [92c] Peter B. Danzig, Deborah Estrin, Jong Suk Ahn, and Brenda Timmerman. "Hybrid Technique for Simulating High Bandwidth-Delay Computer Networks." Extended abstract, p. 260-261, ACM SIGMETRICS 1993.

[91f] Peter Danzig and Deborah Estrin. "Techniques for large internet simulation." USC Network and Distributed Systems Laboratory TR 91-02, November 1991.

[91g] Peter B. Danzig and Sugih Jamin. "tcplib: A library of TCP/IP internetwork traffic characteristics". USC Network and Distributed Systems Laboratory TR91-01, October 1991.

[89c] Peter B. Danzig. "Optimally Selecting the Parameters of adaptive backoff algorithms for computer networks and multiprocessors. Ph.D. Dissertation. U.C. Berkeley, December 1989.